

Discourse Preferences in Dynamic Logic*

Jan Jaspars
CWI
P.O. Box 94079,
1090 GB Amsterdam,
The Netherlands
jaspars@cwi.nl

Megumi Kameyama
SRI International
333 Ravenswood Ave,
Menlo Park,
CA 94025, U.S.A.
megumi@ai.sri.com

Abstract

In order to enrich dynamic semantic theories with a ‘pragmatic’ capacity, we combine dynamic and nonmonotonic (preferential) logics in a modal logic setting. We extend a fragment of Van Benthem and De Rijke’s dynamic modal logic with additional preferential operators in the underlying static logic, which enables us to define defeasible (pragmatic) entailments over a given piece of discourse. We will show how this setting can be used for a dynamic logical analysis of preferential resolutions of ambiguous pronouns in discourse.

1 Introduction

The goal of model-theoretic semantics is to establish an interpretation function from the expressions of a given language to a class of well-understood mathematical structures (models). This enables a formal logical understanding of what an expression means and what its consequences are. For instance, natural language semantics has recently developed a relatively simple *dynamic* model-theoretic understanding of the interplay between indefinite descriptions and anaphoric bindings. These dynamic semantic theories of natural language give model-theoretic explanations of *possible* anaphoric bindings, assuming that additional pragmatics will address the issues of anaphora resolution. A correct dynamic semantic analysis predicts each of the possible referents available in the context, just as a classical logical analysis ‘lists’ all possible scoping and lexical ambiguities.

Consider the following simple discourses (1) and (2).

(1) John met Bill at the station. He_1 greeted him_1 .

(2) Bill met John at the station. He_1 greeted him_1 .

*The first author’s work was supported by CEC project LRE-62-051 (FraCaS). The second author’s work was in part supported by the National Science Foundation and the Advanced Research Projects Agency under Grant IRI-9314961 (Integrated Techniques for Generation and Interpretation). We would like to thank the two anonymous reviewers for helpful comments on an earlier version of the paper.

The two discourses are semantically equivalent. A precise dynamic semantic analysis would treat he_1 and him_1 in both examples as variables that range over the semantic values of *John* and *Bill*, with the additional constraint that the referents of he_1 and him_1 are different. This analysis predicts two sets of equally possible bindings. There is, however, a clear preferential difference between the two discourses. There is a preference for the bindings, $he_1 = \text{John}$ and $him_1 = \text{Bill}$, in (1), and for the opposite bindings, $he_1 = \text{Bill}$ and $him_1 = \text{John}$, in (2).

Preferential effects on discourse interpretations and the entire issue of ambiguity resolution have traditionally been put outside the scope of logical semantics, into the more or less disjoint subfield of ‘pragmatics.’ This academic focus sharply contrasts with the importance placed on disambiguation and resolution issues in natural language processing (or computational linguistics), where realistic accounts of naturally occurring discourses and dialogues are demanded from application systems. Computational accounts, however, often fall short of logical or model-theoretic formalizations. In artificial intelligence (AI), in contrast, logical formalization of pragmatics, or defeasible reasoning, was brought into the central focus of research at an early stage (McCarthy and Hayes, 1969), and led to the development of nonmonotonic logics.

More recently, there are proposals to incorporate defeasible reasoning within natural language semantics to approximate the class of realistic conclusions of a given sentence or discourse (Veltman, 1991; Lascarides and Asher, 1993). In contrast with these specific proposals,¹ we will propose a *general* framework for preferential dynamic semantics, and illustrate how the basic properties of discourse pragmatics exhibited by ambiguous pronouns can be encoded within the framework.

The present framework combines a general model of nonmonotonic logic (Shoham, 1988) and a general model of dynamic logic (Bentham, 1991; de Rijke, 1992). In this logical setup, we specify defeasible information and associated entailment relations over a given discourse, and classify the relative stability of conclusions made on the basis of this additional defeasible information. Our paper is about a general framework of preferential dynamic semantics that abstracts away from numerous specific possibilities for how to represent utterance logical forms and discourse contexts, and how to actually compute preferences. Since logical formalization of discourse pragmatics is in an early stage of development, we believe that it benefits immensely from an attempt such as here to sort out general meta-theoretical issues from specific accounts.

The paper is organized as follows. Section 2 summarizes the preferential effects on ambiguous discourse anaphoric pronouns. Section 3 presents our basic logical framework. Section 4 illustrates formalisms at work in pronoun interpretation in a first-order discourse logic.

2 Preferences in Ambiguous Pronouns

We summarize, here, the basic properties of preferential effects on discourse semantics. We focus on ambiguous pronouns in simple discourses, and illustrate the properties of dynam-

¹Veltman (1991) defines default reasoning in terms of his update semantics. Lascarides and Asher (1993) extend Discourse Representation Theory (DRT) with the definition of commonsense entailment given by Asher and Morreau (1991).

icity, indeterminacy, defeasibility, and preference class interactions.

2.1 Discourse Pragmatics as Preferential Reasoning

Most present-day linguistic theorists assume the trichotomy of syntax, semantics, and pragmatics, but there is no single agreed-upon definition of exactly what *linguistic pragmatics* is. Some equate it with ‘indexicality’, some with ‘context dependence’, and others with ‘language use’ (Levinson, 1983). There is also a common pipeline view of the trichotomy, in that pragmatics adds interpretations to the output of semantics that interprets the output of syntax. In this pipeline view, the direct link between syntax and pragmatics is lost.

We take a logic-inspired definition of pragmatics as the *nonmonotonic* subsystem characterized by *defeasible* rules. We also view all defeasible rules to be *preferences*, so the pragmatics subsystem corresponds to a subspace of preferential reasoning, which *controls* the subspace of *possible* interpretations carved out by the indefeasible linguistic rules in the ‘grammar’ subsystem.² From this perspective, pragmatics is not an underdeveloped subcomponent of semantics alone, but a system that combines all the preferential aspects of phonology, morphology, syntax, semantics, and epistemics. There is evidence that these heterogeneous linguistic preferences interact with one another, and also with nonlinguistic preferences coming from the commonsense world knowledge. What we have then is a dichotomy of grammar and pragmatics subsystems rather than a trichotomy. Under this view, neither indexicality nor context dependence defines pragmatics since there are both indefeasible and defeasible indexical and context-dependent rules. In fact, in a *dynamic* architecture for discourse semantics, where meaning is given to a sequence of sentences rather than to a sentence in isolation, context dependence is an inherent architectural property supporting the anaphoricity of natural language expressions.

2.2 Basic Properties of Discourse Preferences

We will now motivate four basic properties of discourse preferences with examples of ambiguous discourses with ambiguous pronouns. Kameyama (1996) analyzed a survey result of pronoun interpretation preferences from the perspective of interacting preference classes in a dynamic discourse processing architecture. This analysis identified a set of basic ‘design features’ that characterize the preferential effects on discourse meaning, and outlined how they combine to settle on preferred discourse interpretations. These basic properties can be summarized as *dynamicity*, *(in)determinacy*, *defeasibility*, and *preference class interactions*.

Table 1 shows those examples discussed by Kameyama (1996). In a survey, speakers had to pick the preferred reference of pronouns in the last sentence of each discourse example (shown in *italics*).³ Table 2 shows the survey results.⁴ These and similar examples will be

²We assume, following the theoretical linguistic tradition, that there is a linguistic rule system consisting of indefeasible rules of morphosyntax and semantics, and call it the ‘grammar subsystem’. We also assume that most commonsense rules are defeasible, but leave the question open as to whether there are also indefeasible commonsense rules.

³The respondents were told to read the discourses with a ‘neutral’ intonation, for the survey was intended to investigate only *unstressed* pronouns.

⁴The $\chi^2_{df=1}$ significance for each example was computed by adding an evenly divided number of the ‘unclear’ answers to each explicitly selected answer, reflecting the assumption that an ‘unclear’ answer

Grammatical Effects:	
A.	John hit Bill. Mary told <i>him</i> to go home.
B.	Bill was hit by John. Mary told <i>him</i> to go home.
C.	John hit Bill. Mary hit <i>him</i> too.
D.	John hit Bill. <i>He</i> doesn't like <i>him</i> .
E.	John hit Bill. <i>He</i> hit <i>him</i> back.
K.	Babar went to a bakery. He greeted the baker. <i>He</i> pointed at a blueberry pie.
L.	Babar went to a bakery. The baker greeted him. <i>He</i> pointed at a blueberry pie.
Commonsense Effects:	
F.	John hit Bill. <i>He</i> was severely injured.
G.	John hit Arnold Schwarzenegger. <i>He</i> was severely injured.
H.	John hit the Terminator. <i>He</i> was severely injured.
I.	Tommy came into the classroom. He saw Billy at the door. He hit him on the chin. <i>He</i> was severely injured.
J.	Tommy came into the classroom. He saw a group of boys at the door. He hit one of them on the chin. <i>He</i> was severely injured.

Table 1: Discourse Examples in the Survey

used in this paper.

2.2.1 Dynamicity

We are interested in discourse pragmatics, that is, discourse semantics enriched with preferences, so it is natural to start from where discourse semantics leaves off, not losing what discourse semantics has accomplished with its dynamic architecture and the view of sentence meaning as its context change potential. We thus take *dynamicity* to be a basic architectural requirement in an integrated theory of discourse semantics and pragmatics.⁵

The discourse examples (1) and (2), repeated here, demonstrate the fact that the preferred interpretation of an utterance depends on the preceding discourse context.

(1) John met Bill at the station. He_1 greeted him_1 .

(2) Bill met John at the station. He_1 greeted him_1 .

The two discourses are semantically equivalent. Two male persons, ‘John’ and ‘Bill’, engage themselves in a symmetric action of **meeting**. Both individuals are available for anaphoric reference in the next sentence, and since the two pronouns in *He greeted him* must be disjoint in reference and each pronoun has two possible values, dynamic semantic theories predict two equally possible interpretations, **John greeted Bill** and **Bill greeted John**. However, these discourses have different *preferred values* for these pronouns. In (1), due to

shows a genuine ambiguity.

⁵There are two levels of dynamicity that affect utterance interpretation in discourse. One is the utterance-by-utterance dynamicity that affects the overall discourse meaning, and the other is the word-by-word or constituent-by-constituent dynamicity that affects the meaning of the utterance being interpreted. In this paper, we will focus on the former.

Answers				$\chi^2_{df=1}$	p
A.	John 42	Bill 0	Unclear 5	37.53	$p < .001$
B.	John 7	Bill 33	Unclear 7	14.38	$p < .001$
C.	John 0	Bill 47	Unclear 0	47	$p < .001$
D.	J. dislikes B. 42	B. dislikes J. 0	Unclear 5	37.53	$p < .001$
E.	John hit Bill 2	Bill hit John 45	Unclear 0	39.34	$p < .001$
K.	Babar 13	Baker 0	Unclear 0	13	$p < .001$
L.	Babar 3	Baker 10	Unclear 0	3.77	$.05 < p < .10$
F.	John 0	Bill 46	Unclear 1	45.02	$p < .001$
G.	John 24	Arnold 13	Unclear 10	2.57	$.10 < p < .20$
H.	John 34	Terminator 6	Unclear 7	16.68	$p < .001$
I.	Tommy 3	Billy 17	Unclear 1	9.33	$.001 < p < .01$
J.	Tommy 10	Boy 7	Unclear 3	0.45	$.50 < p < .70$

Table 2: Survey Results

a *grammatical parallelism preference* (exhibited by discourse D in Table 1), the preferred interpretation is **John greeted Bill**. In (2), the same parallelism preference leads to the reverse interpretation of **Bill greeted John**.

Dynamic semantics has been motivated by examples such as *A man walks in the park. He whistles.*, where an existential scope extends beyond the syntactic sentence boundary to bind pronouns. Analogously, preferential dynamic semantics would have to account for examples such as (1) and (2), where different syntactic configurations of the same semantic content have different *extended effects* on the preferred interpretation of pronouns.

2.2.2 (In)determinacy

One notable feature of the survey results shown in Table 2 is that the resulting $\chi^2_{df=1}$ significance varies widely. We consider preference to be *significant* if $p < .05$, *weakly significant* if $.05 < p < .10$, and *insignificant* if $.10 < p$ as a straightforward application of elementary statistics. It is reasonable to assume that the statistical significance of a preference corresponds to how determinate the given preference is. Significant preferences are thus unambiguous and determinate, and insignificant preferences indicate ambiguities and indeterminacies. The preferential machinery then must allow both unambiguous and ambiguous preferences to be concluded, rather than always producing a single maximally preferred conclusion.

Preferential reasoning is supposed to resolve ambiguities, however, and unresolved preferential ambiguities make discourses incoherent. It seems reasonable to assume a discourse pragmatic meta-principle that says, *a discourse should produce a single maximally preferred interpretation*. Such a meta-principle is akin to Grice-style maxims of conversation, where a preferred discourse is truthful, adequately informative, perspicuous, relevant, and so forth (Grice, 1975). It seems that this kind of a meta-principle is needed to assure that speakers try to avoid indeterminate preferences precisely because the underlying preferential logical machinery does not guarantee determinacy.

We thus identify a basic property of preferential reasoning — preferential conclusions are

sometimes *determinate* with a single maximally preferred interpretation, and other times *indeterminate* with multiple maximally preferred interpretations. The latter results in a genuine ambiguity, or incoherence, violating the basic pragmatic felicity condition.

Let us turn to concrete examples. Both discourses (1) and (2) have determinate preferred interpretations due to the grammatical parallelism preference. In contrast, discourse (3) leads to no clear preference because no relevant preferences converge on a single determinate choice. Discourse (3) is thus infelicitous.

(3) John and Bill met at the station. He greeted him.

2.2.3 Defeasibility

A conclusion is *defeasible* if it may have to be retracted when some additional facts are introduced. This property is also called *nonmonotonicity*, and is the defining property of *preferences*. This property also defines *pragmatic*, as opposed to grammatical, conclusions under the present assumption that grammatical conclusions are indefeasible.

The following continuation of (1) illustrates defeasibility.

(4) John met Bill at the station. He greeted him. John greeted him back.

In (4), the third sentence, with its indefeasible semantics associated with the adverb *back* (as in discourse E in Table 1), forces a reversal of the preferred interpretation concluded after the second sentence. This on-line reversal produces a discourse-level *garden path* effect, analogous to the sentence-level phenomena such as in *The horse passed the barn fell* or *The astronomer married a star*.

Garden path effects are cases of *preference reversal*, which should not be confused with explicit retractions or negations of indefeasible conclusions. The former can be triggered implicitly, whereas the latter must be explicitly asserted. The latter is illustrated by the following discourse-level *repair* example, where the explicit retraction signal *No* negates the immediately preceding assertion, and opens a way for a different fact to be asserted in the next sentence.

(5) John met Bill at the station. No. He met Paul there.

2.2.4 Preference Classes

When multiple preferences simultaneously succeed, the combined effects are quite unlike the familiar patterns of grammatical rule interactions. When mutually contradictory indefeasible rules both succeed, the whole interpretation is supposed to fail. For instance, *John met Mary at the station. He knows that she loves himself.* leads to no indefeasible interpretation. In contrast, preferences may *override* other preferences that contradict them. Ambiguities persist only when mutually contradictory preferences are equally strong. A logical model of preferential reasoning, therefore, must predict ambiguity resolutions due to overrides.

One type of override is predicted by the so-called Penguin Principle, where the conclusion based on a more specific premise wins (see Lascarides and Asher (1993) for a linguistic application). This principle does not explain all the override phenomena in pragmatic reasoning, however. We must posit the existence of *preference classes* to predict overrides

among groups of preferences (Kameyama, 1996). We thus distinguish between two kinds of conflict resolutions in pragmatics, one due to the Penguin Principle and the other due to preference class overrides.⁶ In this paper, we focus on the interaction between two major preference classes — the *syntactic preferences* based on the *surface structure* of utterances⁷ and the *commonsense preferences* based on the *commonsense world knowledge*.

First consider two examples (A and B) in Table 1 repeated here.

(6) John hit Bill. Mary told him to go home.

(7) Bill was hit by John. Mary told him to go home.

Discourses (6) and (7) illustrate a syntactic preference — the preference for the main grammatical subject to be the antecedent for a pronoun in the next utterance. Henceforth, this syntactic preference is called the *subject antecedent preference*. In (6), the preferred value of the pronoun *him* is John. In (7), with passivization, the preferred value shifts to Bill. Since passivization does not affect the thematic roles (such as Agent or Theme) of these referents, we conclude that this preference shift is directly caused by the shift in grammatical functions.

Next, consider the following.⁸

(8) John hit Bill. He got injured.

(9) The wall was hit by a champagne glass. It broke into pieces.

Discourses (8) and (9) illustrate that the above subject antecedent preference is overridden by a stronger class of preferences having to do with commonsense causal knowledge — in these cases, about hitting causing injuring or breaking.

We thus assume that there are preference classes, or modules, that independently conclude the preferred interpretation of an utterance, and that these class-internal conclusions interact in a certain general overriding pattern to produce the final preference. Table 3 shows the survey result analyzed from this perspective of preference class interactions. Based on this analysis, we will model the following general patterns of preference interactions:

- Indefeasible syntax and semantics override all preferences.
- Commonsense preferences override syntactic preferences.⁹

⁶Asher and Lascarides (1995) implement a class-level override in terms of a ‘meta-penguin principle’ forced on rule classes. Their law of ‘Lexical Impotence’ (p. 96) predicts that discourse inferences generally override default lexical inferences.

⁷This includes both the parallelism and attentional preferences discussed by Kameyama (1996). It was conjectured there that these preference classes may be independent subclasses of a larger ‘entity-level’ preference class, which is qualitatively different from the ‘propositional-level’ commonsense preference class.

⁸(8) is a slight variation of F in Table 1. (9) is a variant of Len Schubert’s (personal communication) example.

⁹This overriding can be difficult when the syntactic preference is extremely strong. For instance, example I in Table 1 creates an utterance-internal garden-path effect where the first syntactically preferred choice for Tommy is retracted in favor of a more plausible interpretation supported by commonsense preferences.

	Syntactic Pref.	Commonsense Pref.	Semantics	Winner
A.	John	unclear	—	Syntactic Pref.
B.	Bill	unclear	—	Syntactic Pref.
C.	John	unclear	Bill	Semantics
D.	John–Bill	unclear	—	Syntactic Pref.
E.	John–Bill	unclear	Bill–John	Semantics
K.	Babar	unclear	—	Syntactic Pref.
L.	Baker	unclear	—	Syntactic Pref.
F.	John	Bill	—	Commonsense Pref.
G.	John	John/Arnold	—	Commonsense Pref.
H.	John	John	—	Commonsense Pref.
I.	Tommy	Billy	—	Commonsense (but difficult)
J.	Tommy	Boy(/Tommy)	—	??

Table 3: Preference Interactions

- Syntactic preferences dominate the final interpretation only if there are no relevant commonsense preferences.

The general overriding pattern we identify here is schematically shown as follows, where \geq represents a ‘can override’ relation:

Indefeasible Syntax and Semantics	\geq	Commonsense Preferences	\geq	‘Syntactic’ Preferences
--------------------------------------	--------	----------------------------	--------	----------------------------

There are a number of questions about these preference classes. For instance, how do they arise, how many classes are there, and why can some classes override others?¹⁰ In this paper, we simply assume the existence of multiple preference classes with predetermined override relationships, and propose a logical machinery that implements their interactions.

We will now turn to the logical machinery that will be used to model pragmatic reasoning with the requisite properties of dynamicity, indeterminacy, defeasibility, and preference class interactions.

3 Dynamic Preferential Reasoning

We have chosen to combine dynamics and preferences in a most general logical setting in order to achieve logical transparency and theoretical independence in the following sense. We hope that the logical simplicity facilitates future meta-logical investigations on the interaction of dynamics and preferential reasoning, and enables applications to a wider variety of preferential (defeasible) phenomena. We will thus combine the most general dynamic

¹⁰Kameyama (1996) proposed that there are three preference classes that respectively concern preferred updates of three data structure components of the dynamic context. These three preference classes also seem to correspond with the three classes of *discourse coherence relations* independently proposed by Kehler (1995) to account for the constraints on ellipsis and other cohesive forms. This indicates a potential integration of two apparently unrelated notions — dynamic context data structure components and coherence relations.

logical approach and the most general logical approach to defeasible reasoning we know. The dynamic (relational) setting consists of the core of the so-called dynamic modal logic of Van Benthem (1991) and de Rijke (1992). Our encoding of defeasibility follows Shoham's (1988) preferential modeling of nonmonotonic logics.

Subsection 3.1 will outline dynamic modal logic, following Jaspars and Krahmer's (1996) fragment of the original logic.¹¹ This part encodes the dynamicity property. Subsections 3.2 and 3.3 will show how preferential reasoning can be accommodated within this fragment of dynamic modal logic. This addition encodes defeasibility, indeterminacy, and differentiation of preference classes. Finally, Subsection 3.4 discusses possible pragmatic meta-constraints on preferential interpretation definable in this logical setting.

3.1 Basic Dynamic Modal Logic

Jaspars and Krahmer (1996) present specifications of current dynamic semantic theories in terms of dynamic modal logic (DML), and show how DML can be used as a universal setting in which the differences and similarities among different dynamic semantic theories can be clarified. The underlying philosophy of this unified dynamics is that dynamic theories evolve from 'dynamifying' an ordinary logic by implementing an order of information growth over the models of this logic.

To start with, one chooses a *static language* \mathcal{L} to reason about the content of *information states* S by means of an *interpretation function*: $\llbracket \cdot \rrbracket : \mathcal{L} \longrightarrow 2^S$. This setting most often consists of a (part of) well-known logic interpreted over a class of well-known models. These models are then taken to be the units of information, that is, information states, within the dynamic modal framework. The second (new) step consists of a definition of an *order of information growth*, \sqsubseteq , over these information states. We write $s \sqsubseteq t$ whenever the state t contains more information than s according to this definition. The conclusive step is the choice of the dynamic language \mathcal{L}^* , which essentially comes down to selecting different dynamic modal operators for reasoning about the relation \sqsubseteq . The triple $\langle S, \sqsubseteq, \llbracket \cdot \rrbracket \rangle$ is also called an \mathcal{L} -*information model*.

Conventions. If $M = \langle S, \sqsubseteq, \llbracket \cdot \rrbracket \rangle$ is an \mathcal{L} -information model, then we write $s \sqsubset t$ whenever $s \sqsubseteq t$ and not $t \sqsubseteq s$. The state t is called a *proper extension* of s . If $T \subseteq S$ then the *minimal states* in T is the set $\{t \in T \mid \forall s \in T : s \sqsubseteq t \Rightarrow t \sqsubseteq s\}$. We will assume that every nonempty subset of information states contains minimal states. Most often, dynamic semantic theories can be described on the basis of information models that satisfy this constraint.

3.1.1 Static and Dynamic Meaning

On the basis of these information models, one can distinguish between static and dynamic meanings of propositions. The *static meaning* of a proposition $\varphi \in \mathcal{L}$ with respect to an \mathcal{L} -information model $M = \langle S, \sqsubseteq, \llbracket \cdot \rrbracket \rangle$, written as $\llbracket \varphi \rrbracket_M$, is the same as $\llbracket \varphi \rrbracket$. The reason is that we want to define a dynamic modal extension \mathcal{L}^* on top of \mathcal{L} , which requires static interpretation as well ($\llbracket \cdot \rrbracket_M : \mathcal{L}^* \longrightarrow 2^S$).

¹¹To be precise, the relational part of this setting is a fragment of the relational expressivity of original dynamic logic.

Given the relational structure, that is, the preorder of information growth \sqsubseteq , over the information states S , we are able to define a *dynamic meaning* of a proposition. Roughly speaking, the dynamic meaning of a proposition is understood as its *effect* on a given information state $s \in S$.¹² In other words, we wish to define the meaning(s) of a proposition φ in the context of an information state $s \in S$: $\llbracket \varphi \rrbracket_{M,s}$.

In general, different dynamic interpretations of a proposition φ are defined according to how φ operates on an information state. For example, φ might be added to or retracted from an information state, or, in a somewhat more complicated case, φ may describe the content of a revision to an information state. Given such an operation o , we will define the o -meaning of a proposition φ with respect to an information state $s \in S$ (in M): $\llbracket \varphi \rrbracket_{M,s}^o$. The proposition φ is the *content* of an operation and o specifies the *type* of operation. In DML, all these operations are defined in terms of the growth relation \sqsubseteq .

Jaspars and Krahmer (1996) postulate that in most well-known logics of mental action or change, we need only four basic operation types: *extension* (+) and *reduction* (−), and their minimal counterparts, *update* (+ μ) and *downdate* (− μ). Given an information order \sqsubseteq for a given set of information states S , these actions are defined as follows:

$$(10) \quad \begin{aligned} \llbracket \varphi \rrbracket_{M,s}^+ &= \{t \in S \mid s \sqsubseteq t, t \in \llbracket \varphi \rrbracket_M\} \\ \llbracket \varphi \rrbracket_{M,s}^- &= \{t \in S \mid t \sqsubseteq s, t \notin \llbracket \varphi \rrbracket_M\} \\ \llbracket \varphi \rrbracket_{M,s}^{+\mu} &= \{t \in \llbracket \varphi \rrbracket_{M,s}^+ \mid \forall u \in S : u \in \llbracket \varphi \rrbracket_{M,s}^+ \ \& \ u \sqsubseteq t \Rightarrow t \sqsubseteq u\} \\ \llbracket \varphi \rrbracket_{M,s}^{-\mu} &= \{t \in \llbracket \varphi \rrbracket_{M,s}^- \mid \forall u \in S : u \in \llbracket \varphi \rrbracket_{M,s}^- \ \& \ t \sqsubseteq u \Rightarrow u \sqsubseteq t\}. \end{aligned}$$

Furthermore, for every action type o we use $\llbracket \varphi \rrbracket_{M,T}^o$ as an abbreviation of the set $\bigcup_{s \in T} \llbracket \varphi \rrbracket_{M,s}^o$ (the o -meaning of φ with respect to T) for all $T \subseteq S$. A special instance of particular importance is the o -meaning with respect to the minimal states in M : $\min_M = \{s \in S \mid \forall t \in S : t \sqsubseteq s \Rightarrow s \sqsubseteq t\}$. We write $\llbracket \varphi \rrbracket_{M,\min}^o$ instead of $\llbracket \varphi \rrbracket_{M,\min_M}^o$, and refer to this set as the minimal o -meaning of φ in M . This is the meaning of a proposition with respect to an empty context. We will also use the notation $\min_M T$ for a given subset $T \subseteq S$ of minimal states in T . We assumed above that $\min_S T \neq \emptyset$ whenever $T \neq \emptyset$, and therefore, $\llbracket \varphi \rrbracket_{M,s}^+ \neq \emptyset \Rightarrow \llbracket \varphi \rrbracket_{M,s}^{+\mu} \neq \emptyset$ (the same holds for $-$ with respect to $-\mu$).

Dynamic semantic theories most often describe relational meanings of propositions obtained from abstractions over the context. For every operation o , we will call the relational interpretation the o -meaning of φ (in M).

$$(11) \quad \llbracket \varphi \rrbracket_M^o = \{\langle s, t \rangle \mid t \in \llbracket \varphi \rrbracket_{M,s}^o\}.$$

Finally, a dynamic modal extension \mathcal{L}^* of \mathcal{L} can be defined. It supplies unary dynamic modal operators of the form $[\varphi]^o$ and $\langle \varphi \rangle^o$, whose static interpretations are as follows:

$$(12) \quad \begin{aligned} \llbracket [\varphi]^o \psi \rrbracket_M &= \{s \in S \mid \llbracket \varphi \rrbracket_{M,s}^o \subseteq \llbracket \psi \rrbracket_M\} \\ \llbracket \langle \varphi \rangle^o \psi \rrbracket_M &= \{s \in S \mid \llbracket \varphi \rrbracket_{M,s}^o \cap \llbracket \psi \rrbracket_M \neq \emptyset\}. \end{aligned}$$

For example, a proposition of the form $[\varphi]^+ \psi$ means that extending the current state with φ necessarily leads to a ψ -state, while $\langle \varphi \rangle^{-\mu} \psi$ means that it is possible to retract φ from

¹²Note that linguistic actions most often affect the mental state of some chosen agents or interpreters, sharply contrasting with physical actions that affect physical situations, as studied in AI for analysis of so-called frame problems, e.g., Shoham (1988).

the current state in a minimal way and end up with the information ψ . In this paper, we will discuss only the extension (+) and update ($+\mu$) meanings of propositions.

Notational conventions. Let C be a set of connectives. Then we write $\mathcal{L} + C$ for the smallest superset of \mathcal{L} closed under the connectives in C . $\mathcal{L} * C$ denotes the smallest superset of \mathcal{L} closed under the connectives appearing in \mathcal{L} and the connectives in C .

3.1.2 Static and Dynamic Entailment

Entailments are defined as relations between sequences of formulae and single formulae. The former contain the *assumptions* and the latter are the *conclusions* of the entailments. To make concise definitions, we also define the static and dynamic meaning of a sequence $\varphi_1, \dots, \varphi_n$, abbreviated as $\vec{\varphi}$, in a dynamic modal language \mathcal{L}^* . Let $M = \langle S, \sqsubseteq, [\cdot] \rangle \in \mathcal{M}_{\mathcal{L}}$, then

$$(13) \quad \llbracket \vec{\varphi} \rrbracket_M = \bigcap_{i=1}^n \llbracket \varphi_i \rrbracket_M \quad \text{and} \quad \llbracket \vec{\varphi} \rrbracket_M^o = \llbracket \varphi_1 \rrbracket_M^o \circ \dots \circ \llbracket \varphi_n \rrbracket_M^o.^{13}$$

The former part defines the static meaning of $\vec{\varphi}$, and the latter part defines the o -meaning of $\vec{\varphi}$. The o -meaning of $\vec{\varphi}$ is the relation of input/output pairs of consecutively o -executing (expanding, updating,...) φ_1 through φ_n .

We will subsequently write $\llbracket \vec{\varphi} \rrbracket_{M,s}^o$ for the set $\{t \in S \mid \langle s, t \rangle \in \llbracket \vec{\varphi} \rrbracket_M^o\}$ and $\llbracket \vec{\varphi} \rrbracket_{M,T}^o = \bigcup_{s \in T} \llbracket \vec{\varphi} \rrbracket_{M,s}^o$ for all $s \in S$ and $T \subseteq S$. We will write $\llbracket \vec{\varphi} \rrbracket_{M,\min}^o$ for the minimal o -meaning of the sequence $\vec{\varphi}$.

Definition 1 Let \mathcal{M} be some class of \mathcal{L} -information models, and let $\varphi_1, \dots, \varphi_n, \psi$ be propositions of some dynamic modal extension \mathcal{L}^* of \mathcal{L} . We define the following entailments for discourse $\varphi_1, \dots, \varphi_n$ ($\vec{\varphi}$):

- $\vec{\varphi}$ *statically entails* ψ with respect to \mathcal{M} if $\llbracket \vec{\varphi} \rrbracket_M \subseteq \llbracket \psi \rrbracket_M$.
- $\vec{\varphi}$ *dynamically entails* ψ according to the operation o (or $\vec{\varphi}$ o -entails ψ) with respect to \mathcal{M} if $\llbracket \vec{\varphi} \rrbracket_{M,s}^o \subseteq \llbracket \psi \rrbracket_M$ for all $M \in \mathcal{M}$ and s in M .
- $\vec{\varphi}$ *minimally o -entails* ψ with respect to \mathcal{M} if $\llbracket \vec{\varphi} \rrbracket_{M,\min}^o \subseteq \llbracket \psi \rrbracket_M$ for all $M \in \mathcal{M}$.

We use $\vec{\varphi} \models_{\mathcal{M}} \psi$, $\vec{\varphi} \models_{\mathcal{M}}^o \psi$ and $\vec{\varphi} \models_{\mathcal{M}}^{\min o} \psi$ as abbreviations for these three entailment relations, respectively.

Note that if the modal operators $[\varphi]^o$ are present within the dynamic modal language \mathcal{L}^* , then the notion of o -entailment in Definition 1 boils down to the static entailment $\models_{\mathcal{M}} [\varphi_1]^o \dots [\varphi_n]^o \psi$.

When we think of operations as updates as in the following sections, the minimal dynamic meaning of a sequence $\varphi_1, \dots, \varphi_n$ is the same as updating the minimal states (the initial context) consecutively with φ_1 through φ_n . This interpretation is the one we will use for *the* interpretation of a discourse or text $\vec{\varphi}$. Of course, as will be the case for most pragmatic

¹³The operation \circ stands for relational composition. For two relations $R_1, R_2 \subseteq S^2$: $R_1 \circ R_2 = \{\langle s, t \rangle \in S^2 \mid \exists u \in S : R_1(s, u) \ \& \ R_2(u, t)\}$.

inferences, the minimal states of an information model should not be states of complete ignorance. To draw the defeasible conclusions discussed in the previous section, we need to add some defeasible background information. For this purpose we need the following notation. If $\Gamma \subseteq \mathcal{L}^*$, then we write \mathcal{M}_Γ for the subclass of models in \mathcal{M} that supports all the formulae in Γ : $\{M = \langle S, \sqsubseteq, \llbracket \cdot \rrbracket \rangle \in \mathcal{M} \mid \llbracket \gamma \rrbracket_M = S \text{ for all } \gamma \in \Gamma\}$. The entailment $\bar{\varphi} \models_{M_\Gamma}^{\min + \mu} \psi$ covers the interpretation of a discourse $\bar{\varphi}$ in the context or background knowledge of Γ .

3.2 Simple Preferential Extensions

Shoham (1988) introduced preferential reasoning into nonmonotonic logics. The central idea is to add a preferential structure over the models of the logic chosen as the inference mechanism. This preferential structure is most often some partial or pre-order. A nonmonotonic inference, $\varphi_1, \dots, \varphi_n \approx \psi$, then says that ψ holds in all the maximally preferred $\bar{\varphi}$ -models.

In many nonmonotonic formalisms such as Reiter's (1980) default logic, an additional preferential structure of an assumption set $\bar{\varphi}$ is specified by explicit *default assumptions* Δ , which are defeasible. The central idea is to use 'as much information from Δ as possible' as long as it is consistent with the strict assumptions Φ . We will also encode this maximality preference in our definition. In this paper, we use a preferential operator \mathbf{p} to specify the additional defeasible information. A proposition of the form $\mathbf{p}\varphi$ refers to the maximally preferred φ -states.

3.2.1 Single Preference Classes

Preferential reasoning can be accommodated within the DML framework by assigning an additional preferential structure to the space of information states. There are essentially two ways to do this. In one method, the preferential structure is added to the static structure over information states ($\llbracket \cdot \rrbracket$), and in the other method, it is added to the dynamic structure on these states (\sqsubseteq). We take the first, simpler, option in this paper.¹⁴

As explained in Subsection 2.2.4, the preferential reasoning for anaphoric resolution needs to take different *preference classes* into account. In Subsection 3.3, we will give DML-style definitions for such structures, which will be a straightforward generalization of the following definition of a single preference class.

Definition 2 Extension with a single preference class:

- A *single preferential extension* $\mathcal{L}_\mathbf{p}$ of the static language \mathcal{L} is the smallest superset of \mathcal{L} such that $\mathbf{p}\varphi \in \mathcal{L}_\mathbf{p}$ for all $\varphi \in \mathcal{L}$.
- A *preferential* \mathcal{L} -model is an information $\mathcal{L}_\mathbf{p}$ -model $M = \langle S, \sqsubseteq, \llbracket \cdot \rrbracket \rangle$, with $\llbracket \cdot \rrbracket$ representing a pair of interpretation functions $\langle {}^0\llbracket \cdot \rrbracket, {}^1\llbracket \cdot \rrbracket \rangle$ such that $M_0 = \langle S, \sqsubseteq, {}^0\llbracket \cdot \rrbracket \rangle$ and $M_1 = \langle S, \sqsubseteq, {}^1\llbracket \cdot \rrbracket \rangle$ are \mathcal{L} -information models, and $\llbracket \varphi \rrbracket = {}^0\llbracket \varphi \rrbracket$ and $\llbracket \mathbf{p}\varphi \rrbracket = {}^1\llbracket \varphi \rrbracket$ for all $\varphi \in \mathcal{L}$.

¹⁴The latter, more complex, option would be a more balanced combination of dynamic and preferential reasoning because the preferential structure is represented at the same level of information order over which dynamicity is defined. From this perspective, the preferential structuring of models of a given logic that supplies a nonmonotonic component is analogous to dynamifying a logic by informational structuring as described by Jaspars and Krahmer (1996). Such investigations are left for a future study.

- If \mathcal{M} is a class of \mathcal{L} -information models, then the class of all preferential \mathcal{L} models whose nonpreferential part (0) is a member of \mathcal{M} is called the single-preferential enrichment of \mathcal{M} .
- If $\mathcal{L}^* = \mathcal{L} + (*) C$, then \mathcal{L}_p^* refers to the language $\mathcal{L}_p + (*) C$.

The interpretation function $\llbracket \cdot \rrbracket$ consists of an *indefeasible part* ${}^0\llbracket \cdot \rrbracket$ and a *defeasible part* ${}^1\llbracket \cdot \rrbracket$. Both parts are interpretation functions of the static language: ${}^{0,1}\llbracket \cdot \rrbracket : \mathcal{L} \longrightarrow \mathbf{2}^S$. The indefeasible part replaces the ordinary interpretation function, while the additional defeasible part is the ‘pragmatic’ strengthening of this standard reading. Note that a preferential extension gives us a set of preferred states, allowing both determinate and indeterminate interpretations.

3.2.2 Dynamic Preferential Meaning and Preferential Entailment

Definition (14) illustrates the static and dynamic *preferential* meaning of a sentence φ analogous to the nonpreferential definitions presented in Subsection 3.1.1. The *static preferential meaning* of a sentence φ (in a model M) is written as $\langle\langle \varphi \rangle\rangle_M$, and the ‘dynamic’ *preferential meaning of φ with respect to a given information state (context) s* in a model M is written as $\langle\langle \varphi \rangle\rangle_{M,s}$.

$$(14) \quad \begin{aligned} \langle\langle \varphi \rangle\rangle_M &= \llbracket \mathbf{p} \varphi \rrbracket (= {}^1\llbracket \varphi \rrbracket) \\ \langle\langle \varphi \rangle\rangle_{M,s} &= \begin{cases} \llbracket \mathbf{p} \varphi \rrbracket_{M,s}^o & \text{if } \llbracket \mathbf{p} \varphi \rrbracket_{M,s}^o \neq \emptyset \\ \llbracket \varphi \rrbracket_{M,s}^o & \text{otherwise.} \end{cases} \end{aligned}$$

In line with the definitions of Subsection 3.1.1, we write $\langle\langle \varphi \rangle\rangle_M^o$ for the relational abstraction of $\langle\langle \varphi \rangle\rangle_{M,s}^o$. Our definition of the preferential dynamic meaning of a discourse $\varphi_1, \dots, \varphi_n = \vec{\varphi}$ is written as $\langle\langle \vec{\varphi} \rangle\rangle_{M,s}^o$, and its definition deviates from the way $\llbracket \vec{\varphi} \rrbracket_{M,s}^o$ has been defined above because a simple relational composition of the preferential dynamic readings of single sentences does not give us a satisfactory definition. The failure of normal composition in this respect can be illustrated by the following simple abstract example. Suppose $\vec{\varphi} = \varphi_1, \varphi_2$ is a two-sentence discourse with

$$(15) \quad \llbracket \mathbf{p} \varphi_1 \rrbracket_{M,a}^{+\mu} = \{b, c\}, \llbracket \varphi_2 \rrbracket_{M,1}^{+\mu} = \{d\}, \llbracket \mathbf{p} \varphi_2 \rrbracket_{M,1}^{+\mu} = \emptyset \text{ and } \llbracket \mathbf{p} \varphi_2 \rrbracket_{M,2}^{+\mu} = \{e\}.$$

We obtain both $\langle a, d \rangle, \langle a, e \rangle \in \langle\langle \varphi_1 \rangle\rangle_M \circ \langle\langle \varphi_2 \rangle\rangle_M$. The second pair $\langle a, e \rangle$ is composed of maximally preferred readings while the first pair $\langle a, d \rangle$ is not. Because these two pairs are both equal members of the composition, such a definition of the preferential meaning of a discourse is not satisfactory.

The two-sentence discourse in this example has four possible readings: (1) composing the two defeasible/preferential readings, (2) composing the indefeasible reading of one sentence and the defeasible reading of the other sentence in two possible orders, and (3) composing the two indefeasible readings. As we said earlier, it is reasonable to use as much preference as possible, which means that (1) should be the ‘best’ composition, the two possibilities in (2) should be the next best, and (3) should be the ‘worst’. We will encode this preferential ordering based on the amount of preferences into the entailment definition. What about then the two possible ways of mixing indefeasible and defeasible readings of the two sentences in

the case of (2)? A purely amount-based comparison would not differentiate them. Are they equally preferred?

In addition to the sensitivity to the amount of overall preferences, we hypothesize that the discourse's linear progression factor also gives rise to a preferential ordering. We thus distinguish between the two compositions of infeasible and defeasible readings in (2), and assign a higher preference to the composition in which the first sentence has the defeasible/preferential reading rather than the infeasible reading. The underlying intuition is that the defeasibility of information is inversely proportional to the flow of time. It is harder to defeat conclusions drawn earlier in the given discourse. This has to do with the fading of nonsemantic memory with time. Earlier (semantic) conclusions tend to persist, while explicit sentence forms fade away as discourse continues. It seems easier to distinguish (defeasible) conclusions from recently given information than from information given earlier.

We thus take the preferential context-sensitive reading of a discourse $\vec{\varphi} = \varphi_1, \dots, \varphi_n$ to be the interpretation that results from applying preferential rules as *much* as possible and as *early* as possible. This type of interpretation can be defined on the basis of an induction on the length of discourses:

$$(16) \quad \begin{aligned} {}^{2k} \llbracket \vec{\varphi} \rrbracket_{M,s}^o &= \llbracket \varphi_n \rrbracket_{M,T} & \text{and} \\ {}^{2k+1} \llbracket \vec{\varphi} \rrbracket_{M,s}^o &= \llbracket \mathbf{p} \varphi_n \rrbracket_{M,T} & \text{with } T = {}^k \llbracket \varphi_1, \dots, \varphi_{n-1} \rrbracket_{M,s}^o. \end{aligned}$$

Note that $k < 2^{n-1}$ in this inductive definition. ${}^0 \llbracket \varphi_1 \rrbracket_{M,s}$ and ${}^1 \llbracket \varphi_1 \rrbracket_{M,s}$ are given by the \mathcal{L} -information model M . The set of states ${}^k \llbracket \vec{\varphi} \rrbracket_{M,s}^o$ is called the *o*-meaning of $\vec{\varphi}$ of *priority* k with respect to s in M . In this way, we obtain 2^n readings of a given discourse. The *preferential o-meaning* of a discourse $\vec{\varphi}$ (w.r.t. s in M) is then the same as the nonempty interpretation of the highest priority larger than 0, and if all these readings are empty, then the preferential *o*-meaning coincides with the completely infeasible reading of priority 0.

$$(17) \quad \begin{aligned} \langle\langle \vec{\varphi} \rangle\rangle_{M,s}^o &= {}^k \llbracket \vec{\varphi} \rrbracket_{M,s}^o \\ \text{with } k &= \max(\{i \mid {}^i \llbracket \vec{\varphi} \rrbracket_{M,s}^o \neq \emptyset, 0 < i < 2^n\} \cup \{0\}). \end{aligned}$$

Note that application of this definition to example (15) yields $\langle\langle \varphi_1, \varphi_2 \rangle\rangle_{M,0}^{+\mu} = \{e\}$. Definition (17) leads to the following succinct definition of *preferential dynamic entailment*:

$$(18) \quad \varphi_1, \dots, \varphi_n \models_{\mathcal{M}}^o \psi \Leftrightarrow \langle\langle \varphi_1, \dots, \varphi_n \rangle\rangle_{M,s}^o \subseteq \llbracket \psi \rrbracket_M \text{ for all } s \text{ in } M, \text{ for all } M \in \mathcal{M}.$$

This definition says that for every input state of a discourse $\vec{\varphi}$, the maximally preferred readings of the discourse always lead to ψ -states. We write $\vec{\varphi} \models_{\mathcal{M}}^{\min o} \psi$ whenever $\langle\langle \vec{\varphi} \rangle\rangle_{M \min}^o \subseteq \llbracket \psi \rrbracket_M$ for all $M \in \mathcal{M}$ (minimal preferential dynamic entailment).

3.3 Multiple Preference Classes

Now we turn to information models of multiple preference classes needed for formalizing the preference interaction in pronoun resolution, as motivated in Section 2. If we assume a linear priority order on these preference classes, then it is not hard to generalize Definition 2

of a single preference class given in Subsection 3.2.1. We will assume such determinate overriding relations among preference classes here.¹⁵

Definition 3 Extension with multiple preference classes:

- A multiple (m) preferential extension $\mathcal{L}_{p,m}$ of \mathcal{L} is the smallest superset of \mathcal{L} such that $p_i \varphi \in \mathcal{L}_{p,m}$ for all $\varphi \in \mathcal{L}$.
- A multiple (m) preferential \mathcal{L} -model is a $\mathcal{L}_{p,m}$ -information model $\langle S, \sqsubseteq, [\![\cdot]\!] \rangle$ such that $[\![\cdot]\!] = \langle {}^0[\![\cdot]\!], \dots, {}^m[\![\cdot]\!] \rangle$ with $M_i = \langle S, \sqsubseteq, {}^i[\![\cdot]\!] \rangle \in \mathcal{M}_{\mathcal{L}}$ for all $i \in \{0, \dots, m\}$, and $[\![\varphi]\!] = {}^0[\![\varphi]\!]$ and $[\![p_i \varphi]\!] = {}^i[\![\varphi]\!]$ for all $\varphi \in \mathcal{L}$ and $i \in \{1, \dots, m\}$.
- The class of m -preferential enrichments of a class of \mathcal{L} information models \mathcal{M} is the class of all preferential \mathcal{L} models whose indefeasible part (0) is a member of \mathcal{M} .

Intuitively, $p_i \varphi$ says that the current state is a preferred state according to the i -th preference class and the content φ . We use a simple generalization of the preferential dynamic meaning given in the previous section for the singular preference setting. For a given discourse $\vec{\varphi} = \varphi_1, \dots, \varphi_n$, we define $(m+1)n$ readings and define their associated priority in the same manner as in (16). Let $k < (m+1)^{n-1}$ and $T = {}^k[\![\varphi_1, \dots, \varphi_{n-1}]\!]_{M,s}^o$. Then ${}^{(m+1)k}[\![\vec{\varphi}]\!]_{M,s}^o = [\![\varphi_n]\!]_{M,T}^o$ and ${}^{(m+1)k+i}[\![\vec{\varphi}]\!]_{M,s}^o = [\![p_i \varphi_n]\!]_{M,T}^o$ for all $i \in \{1, \dots, m\}$. The preferential o -meaning of $\vec{\varphi}$ with respect to a state s in an information model M , $\langle\langle\vec{\varphi}\rangle\rangle_{M,s}^o$, is defined in the same way as for the single preferential case (17): replace 2 with $m+1$.

3.4 Pragmatic Meta-constraints

For most applications, however, this definition is far too general, and we need to regulate the interplay of indefeasible and defeasible interpretations with additional constraints. We discuss some candidates here. Let $M = \langle S, \sqsubseteq, \langle {}^0[\![\cdot]\!], {}^1[\![\cdot]\!] \rangle \rangle$ be a preferential \mathcal{L} -model.

PRINCIPLE 1 (*Realism*) Every preferential φ -state, or $p \varphi$ -state, is a φ -state itself:¹⁶

$${}^1[\![\varphi]\!] \subseteq {}^0[\![\varphi]\!].$$

This principle is perhaps too strict. In some types of defeasible reasoning, we would like to assign preferential meanings to meaningless or ill-formed input, which would give us the robustness to recover from errors. Such robustness can be expressed in terms of a restriction to nonempty indefeasible readings as follows: ${}^0[\![\varphi]\!] \neq \emptyset \Rightarrow {}^1[\![\varphi]\!] \subseteq {}^0[\![\varphi]\!]$ (Robust Realism).

PRINCIPLE 2 (*Minimal Preference*) In minimal information states, if a proposition has an indefeasible reading, it should also have a preferential reading:

$$[\![\varphi]\!]_{M,\min}^o \neq \emptyset \Rightarrow [p \varphi]_{M,\min}^o \neq \emptyset.$$

¹⁵Kameyama (1996) points out that this is not always the case, but in most cases, strict linearity can be enforced through ‘uniting’ multiple preference classes of an equal strength into a single one: $[(p \cup p')\varphi] = [p\varphi] \cup [p'\varphi]$.

¹⁶Compare with the ‘realism’ principle (Cohen and Levesque, 1990): all intended or goal worlds of a rational agent should be epistemically possible. This constraint is often used to distinguish between an agent’s desires and intentions.

The intuition here is that in a minimal state there should be no obstacles that prevent the interpreter from using his preferential expectations or prejudices. In section 4, we will discuss some variants of this principle, which are required to account for certain anaphora resolution preferences.

PRINCIPLE 3 (*Preservation of Equivalence*) Two propositions with the same indefeasible content should also have the same defeasible content:

$${}^0\llbracket\varphi\rrbracket = {}^0\llbracket\psi\rrbracket \Rightarrow {}^1\llbracket\varphi\rrbracket = {}^1\llbracket\psi\rrbracket.$$

This principle is not always desirable.¹⁷ For example, in discourses (1) and (2), *John met Bill* and *Bill met John* have the same semantic/indefeasible content, but different pragmatic/defeasible readings. However, some weaker types of equivalence preservation need to play a role for a satisfactory treatment of anaphoric resolution. Such weakenings will also be discussed in section 4.

PRINCIPLE 4 (*Complete Determinacy*) Every preferential φ -extension of a given information state s has at most one maximal element.

$$\#({}^1\llbracket\varphi\rrbracket_{M,s}^{+\mu}) \leq 1 \text{ for all } s.$$

This excludes indeterminacy described in Subsection 2.2.2, prohibiting Nixon Diamonds. Intuitively, it says that pragmatics always enforces certainty. In other words, in cases of semantic uncertainty, pragmatics always enforces a single choice. For example, discourse (3) should always lead to a single pragmatic solution. Therefore, as argued earlier, this constraint is also unrealistic.

4 Toward a Preferential Discourse Logic

We will discuss, here, two different instances of preferential extensions of the DML-setting of the previous section. As we have seen, such an instantiation requires a specification of static and dynamic modal languages and a class of information models. In Subsection 4.1, we will discuss a simple propositional logic, and explain how simple defeasible (preferential) propositional entailments can or cannot be drawn from a set of preferential rules. Our examples will illustrate the defeasible inference patterns commonly called the Penguin Principle and the Nixon Diamond. In Subsection 4.2, we will define a much richer dynamic semantics that integrates the defeasible propositional inferences explained in Subsection 4.1 into anaphora resolution preferences. Such a combination is needed to account for the preferential effects on anaphora resolution. In Subsection 4.3, we will define first-order variants of pragmatic meta-constraints. In Subsection 4.4, we will illustrate the first-order preferential discourse logic with discourse examples with ambiguous pronouns as discussed in Section 2.

¹⁷This principle is often used in nonmonotonic logics. It implies, for example, the dominance of the default conclusions from more specific information (${}^0\llbracket\varphi\rrbracket \subseteq {}^0\llbracket\psi\rrbracket \Rightarrow {}^1\llbracket\varphi \wedge \psi\rrbracket = {}^1\llbracket\varphi\rrbracket$). If *penguin* \wedge *bird* is equivalent to *penguin*, then Principle 3 makes all the preferential information based on *penguin* applicable, while the preferential information based on *bird* may be invalid for *penguin* \wedge *bird*.

Static Language (\mathcal{L}):	A set of <i>literals</i> : $\mathcal{P} \cup \{\neg p \mid p \in \mathcal{P}\}$
Dynamic Language (\mathcal{L}^*):	$\mathcal{L} + \{[\cdot]^{+\mu}, \langle \cdot \rangle^{+\mu}\}$
States (S):	arbitrary nonempty set.
Order (\sqsubseteq):	arbitrary preorder over S .
Interpretation ($\llbracket \cdot \rrbracket$):	A function $\mathcal{L} \mapsto \wp(S)$ such that <ul style="list-style-type: none"> (i) $\forall \varphi \in \mathcal{L}, s, t \in S : s \sqsubseteq t, s \in \llbracket \varphi \rrbracket \Rightarrow t \in \llbracket \varphi \rrbracket$. (ii) $\forall p \in \mathcal{P} : \llbracket p \rrbracket \cap \llbracket \neg p \rrbracket = \emptyset$. (iii) $\forall \varphi \in \mathcal{L} : \llbracket \varphi \rrbracket \cap \min_M S = \emptyset$.

Table 4: A Class of Propositional Information Models

4.1 A Simple Propositional Preferential Dynamic Logic

Table 4 gives a DML-specification of a simple dynamic propositional logic. The single preferential extension of this logic illustrates how preferential entailments are established according to the definitions given in the previous section. The information states of this model are partial truth value assignments for the propositional atoms: an atom is either true, false, or undefined. The information order is arbitrary, while the interpretation function is (i) *monotonic*, that is, expansions contain more atomic information and (ii) *coherent*, that is, expansions contain no contradictory information, and furthermore, there is a constraint that (iii) the minimal states have empty atomic content.

Let \mathcal{M} be the class of all single-preferential enrichments of this class of information models subject to both Principle 1 (Realism) and Principle 2 (Minimal Preference) defined in the previous section. Let $\{\text{bird}, \text{penguin}, \text{can-fly}\} \subseteq \mathcal{P}$, and let Γ be the following set of \mathcal{L}_p^* -formulae:

$$(19) \{[\text{p bird}]^{+\mu} \text{can-fly}, [\text{p bird}]^{+\mu} \neg \text{penguin}, [\text{p penguin}]^{+\mu} \neg \text{can-fly}, [\text{penguin}]^{+\mu} \text{bird}\},^{18}$$

then:

$$(20) \quad \text{bird} \approx_{\mathcal{M}_\Gamma}^{\min + \mu} \text{can-fly} \quad \text{and} \quad \text{bird}, \text{penguin} \approx_{\mathcal{M}_\Gamma}^{\min + \mu} \neg \text{can-fly}.$$

This entailment is validated by the following derivation for all models $N \in \mathcal{M}_\Gamma$:

$$\begin{aligned} \langle \text{bird} \rangle_{N, \min}^{+\mu} &= {}^1 \llbracket \text{bird} \rrbracket_{N, \min}^{+\mu} \text{ and} \\ \langle \text{bird}, \text{penguin} \rangle_{N, \min}^{+\mu} &= {}^1 \llbracket \text{bird}, \text{penguin} \rrbracket_{N, \min}^{+\mu} = \\ \llbracket \text{bird}, \text{p penguin} \rrbracket_{N, \min}^{+\mu} &= \llbracket \text{p penguin} \rrbracket_{N, \min}^{+\mu} \subseteq \llbracket \neg \text{can-fly} \rrbracket_N. \end{aligned}$$

By definition of the entailment $\approx_{\mathcal{M}_\Gamma}^{\min + \mu}$, we obtain the results of (20).

Next, suppose that $\{\text{republican}, \text{pacifist}, \text{quaker}\} \subseteq \mathcal{P}$, and

$$(21) \quad \Delta = \{[\text{p quaker}]^{+\mu} \text{pacifist}, [\text{p republican}]^{+\mu} \neg \text{pacifist}\}.$$

Here, the preferential readings of *quaker* and *republican* contradict each other. One may expect that we get $\text{quaker}, \text{republican} \approx_{\mathcal{M}_\Delta}^{\min + \mu} \text{pacifist}$, because the preferences of the last sentence are taken to be weaker in the definition (17). This is not the case, however, because

¹⁸The set Γ prescribes that ‘normal birds can fly’, ‘normal birds are not penguins’, ‘normal penguins cannot fly’ and that ‘penguins are birds’.

it is possible that a *republican* cannot be a normal *quaker* ($[republican]^{+\mu} [\mathbf{p} quaker]^{+\mu} \perp$) or vice versa ($[quaker]^{+\mu} [\mathbf{p} republican]^{+\mu} \perp$).

If such preferential blocks are removed, we obtain order-sensitive entailments:

$$(22) \quad \begin{aligned} quaker, republican &\approx_{\mathcal{M}_{\Delta'}}^{\min+\mu} pacifist && \text{and} \\ republican, quaker &\approx_{\mathcal{M}_{\Delta'}}^{\min+\mu} \neg pacifist \end{aligned} ,$$

with Δ' denoting the set:

$$(23) \quad \Delta \cup \{[quaker]^{+\mu} \langle \mathbf{p} republican \rangle^{+\mu} \top, [republican]^{+\mu} \langle \mathbf{p} quaker \rangle^{+\mu} \top\}.$$

Let \mathcal{N} be the class of double-preferential enrichments of the model given in Table 4 subject to the realism and minimal preference principles on both classes. Let Δ'' be the set

$$(24) \quad \begin{aligned} &\{[\mathbf{p}_1 quaker]^{+\mu} pacifist, [\mathbf{p}_2 quaker]^{+\mu} quaker, \\ &[\mathbf{p}_2 republican]^{+\mu} \neg pacifist\} \quad \cup \\ &\{[quaker]^{+\mu} \langle \mathbf{p}_i republican \rangle^{+\mu} \top, \\ &[republican]^{+\mu} \langle \mathbf{p}_i quaker \rangle^{+\mu} \top \mid i = 1, 2\} \end{aligned} .$$

The second rule says that the \mathbf{p}_2 -reading of *quaker* does not entail any information in addition to its indefeasible reading. In this setting, the two variants in (22) yield the same conclusion dominated by the \mathbf{p}_2 -reading of *republican*:

$$(25) \quad \begin{aligned} quaker, republican &\approx_{\mathcal{M}_{\Delta''}}^{\min+\mu} \neg pacifist && \text{and} \\ republican, quaker &\approx_{\mathcal{M}_{\Delta''}}^{\min+\mu} \neg pacifist \end{aligned} .$$

4.2 A First-order Preferential Dynamic Semantics

We will now come to an analysis of the discourses with ambiguous pronouns discussed in Section 2. Typical dynamic semantic analyses of discourse, such as the relational semantics for dynamic predicate logic (Groenendijk and Stokhof, 1991) or first-order DRT such as presented, for example, in Muskens, Benthem, and Visser (1997),²⁰ do not yield a satisfactory preferential dynamic semantics when we integrate them with the preferential machinery of the previous section. In these types of semantic theories, dynamicity is restricted to the value assignment of variables for interpretation of possible anaphoric links, but to account for anaphora resolution we need a logic that supports a preferential interplay of variable assignments, predicates, names, and propositions. In the terminology of Jaspars and Krahmer (1996), we need to ‘dynamify’ more parameters of first-order logic than just the variable assignments.²¹ To arrive at such extended dynamics over first-order models, we will estab-

¹⁹Take $\top = [p]^{+\mu} p$.

²⁰Jaspars and Krahmer (1996) discuss the DML-specification of this semantics for DRT. On the basis of these DML-specifications, one can transfer the present definitions of preferential dynamic entailment to a range of dynamic semantics.

²¹Benthem and Cepparello (1994) discuss such further dynamification. Groenendijk, Stokhof, and Veltman (1996) propose a semantic theory that combines ‘propositional’ and ‘variable’ dynamics, introducing a dynamic semantics over assignment-world pairs. It may be possible to obtain a suitable preferential extension of this type of semantics for our purposes as well.

States (S):	A collection of quadruples $s = \langle D^s, I_p^s, I_c^s, I_v^s \rangle$ with D^s a nonempty set of <i>individuals</i> , $I_p^s : \text{Pred}^n \rightarrow \wp((D^s)^n)$ the <i>local interpretation of predicates</i> , $I_c^s : \text{Con} \rightsquigarrow D$ a <i>partial local interpretation of constants</i> , and $I_v^s : \text{Var} \rightsquigarrow D$ a <i>partial variable assignment</i> .
Order (\sqsubseteq):	A preorder over S such that <ul style="list-style-type: none"> (i) For all $s, t \in S$ if $s \sqsubseteq t$ then $D^s \subseteq D^t$, $I_p^s(P) \subseteq I_p^t(P)$ for all predicates P, $I_c^s(c) = I_c^t(c)$ for all $c \in \text{Dom}(I_c^s)$ and $I_v^s(x) = I_v^t(x)$ for all $x \in \text{Dom}(I_v^s)$. (ii) For all $s, t \in S$ if $s \sqsubseteq t$, $d \in D^t$ and $x \in \text{Var} \setminus \text{Dom}(I_v^s)$, then there exists $u \in S$ such that $s \sqsubseteq u$ and $D^t = D^u$, $I_p^t = I_p^u$, $I_c^t = I_c^u$, $\text{Dom}(I_v^u) = \text{Dom}(I_v^s) \cup \{x\}$ and $I_v^u(x) = d$. (iii) For all $s \in \min_M S$: $I_p^s(P) = \emptyset$ for all predicates P and $\text{Dom}(I_c^s) = \text{Dom}(I_v^s) = \emptyset$.
Interpretation ($\llbracket \cdot \rrbracket$):	$\llbracket Pt_1 \dots t_n \rrbracket = \{s \in S \mid \langle I_t^s(t_1), \dots, I_t^s(t_n) \rangle \in I_p^s(P)\},$ $\llbracket t_1 = t_2 \rrbracket = \{s \in S \mid I_t^s(t_1) = I_t^s(t_2)\},$ $\llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket, \llbracket \varphi \vee \psi \rrbracket = \llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket, \llbracket \perp \rrbracket = \emptyset.$

Table 5: A Class of First-order Information Models

lish a combination of the ‘ordinary’ dynamics-over-assignments semantics with the models of information growth used in possible world semantics for classes of constructive logics.²²

Let us first present the class of our information models. The basic linguistic ingredients are the same as for first-order logic: Con a set of constants, Var a disjoint countably infinite set of variables, and for each natural number n a set of n -ary predicates Pred^n . The static language is the same as for first-order logic except for quantifiers and negation. The dynamic language supplies the formalism with dynamic modal operators $[\cdot]^{+\mu}$ and $\langle \cdot \rangle^{+\mu}$:

$$\begin{aligned}
\text{Atoms} &= \{Pt_1 \dots t_n \mid P \in \text{Pred}^n, t_i \in \text{Con} \cup \text{Var}\} \\
&\quad \cup \{t_1 = t_2 \mid t_i \in \text{Con} \cup \text{Var}\} \\
\mathcal{L} &= \text{Atoms} + \{\wedge, \vee, \perp\} \\
\mathcal{L}^* &= \mathcal{L} * \{[\cdot]^{+\mu}, \langle \cdot \rangle^{+\mu}\}.
\end{aligned}
\tag{26}$$

Table 5 presents the intended \mathcal{L} -information models. The growth of the information order \sqsubseteq is subject to three constraints. The first one (i) says that all the parameters of first-order logic, that is, the domains, interpretation of predicates and constants, and the variable assignments, grow with the information order. The other two constraints seem unorthodox. The second constraint (ii) ensures the freedom of variables in this setting. It tells us that in each state the range of a ‘fresh’ variable is unlimited, that is, it may have the value of each current or ‘future’ individual. This means that for every individual d in an extension t , every variable x that does not yet have an assigned value may be assigned the value d in a state containing the same information as t . This constraint differentiates the roles of constants and variables in this setting. The last constraint (iii) says that the minimal information states do not contain atomic information. It was also used for propositional information models in Subsection 4.1.

The interpretation function is more or less standard. Verification of an atomic sentence requires determination of all the present terms, also for identities.

²²See Troelstra and Van Dalen (1988) or Fitting (1969) for the case of intuitionistic logic.

Quantification can be defined by means of the dynamic modal operators. For example, (27) means that the *Meet*-relation is symmetric and *Greet*-relation is irreflexive.

$$(27) \quad [\text{Meet } xy]^{+\mu} \text{Meet } yx \text{ and } [\text{Greet } xy]^{+\mu} [x = y]^{+\mu} \perp.$$

Ordinary universal quantification can be defined by using identity and extension modality: $\forall x \varphi = [x = x]^+ \varphi$.²³ Negation can also be defined by means of a dynamic modal operator: $\neg \varphi = [\varphi]^+ \perp$.²⁴ A typical (singular) preferential sentence would be

$$(28) \quad [\mathbf{p} \text{Meet } xy]^{+\mu} [\mathbf{p} \text{Greet } uv]^{+\mu} (u = x \wedge v = y),$$

which means that the concatenation of the preferential reading of a *Meeting* and a *Greeting* pair makes the variables match according to the grammatical parallelism preference.²⁵

4.3 First-order Constraints for Preferential Dynamic Reasoning

To model the preferential effects on ambiguous pronouns discussed in Section 2, we need to postulate several first-order variants of the pragmatic meta-constraints discussed in Subsection 3.4. The first-order expressivity of the languages \mathcal{L} and \mathcal{L}^* given in (26) and the fine structure of the information models presented in Table 5 enable us to calibrate these meta-constraints for preferential interpretation on first-order discourse representations.

We will adopt only Principle 1 (Realism) in its purely propositional form. Three other constraints that we will impose on preferential interpretation regulate some ‘harmless’ interplay of preferences and terms. Let $M = \langle S, \sqsubseteq, \llbracket \cdot \rrbracket \rangle$ be a preferential \mathcal{L} -model with $\llbracket \cdot \rrbracket = \langle {}^0 \llbracket \cdot \rrbracket, {}^1 \llbracket \cdot \rrbracket \rangle$.

To begin with, fresh variables have no content, and therefore, we do not allow them to block preferential interpretation. In other words, a proposition that contains only fresh variables as terms always has a preferential $+\mu$ -reading whenever it has an indefeasible $+\mu$ -meaning. In fact, this is a variant of Principle 2, the principle of minimal preference.

PRINCIPLE 5 (*Minimal Preference for Fresh Variables*) Let s be an information state in an information model of the type described in Table 5. If $\text{Dom}(I_v^s)$ has an empty intersection with the variables occurring in a given proposition φ , and no constants occur in φ , then

$$\llbracket \varphi \rrbracket_{M,s}^{+\mu} \neq \emptyset \Rightarrow \llbracket \mathbf{p} \varphi \rrbracket_{M,s}^{+\mu} \neq \emptyset.$$

The two other constraints for first-order discourses are obtained by weakening Principle 3 (Preservation of Equivalence). Although this principle itself is too strong, we would like to have some innocent logical transparency of the preferential operator. We thus postulate Principles 6 and 7.

²³Note that to get the proper universal reading here, we need to be sure that x is a fresh variable (e.g., in the minimal states).

²⁴A proper definition of existential quantification does not seem feasible since $\langle x = x \rangle^{+\mu} \varphi$ is not persistent. A better candidate is $\neg \forall x \neg \varphi$, which behaves persistently. For \perp we may take $\langle x = x \rangle^+ (x = x)$.

²⁵A general implementation of the parallelism preference would require a second-order scheme.

PRINCIPLE 6 (*Preservation under Renaming Fresh Variables.*) Preferential readings should be maintained when fresh variables are replaced by other fresh variables:

$$\forall x, y \in \mathbf{Var} \setminus \mathbf{Dom}(I_v^s) : s \in \llbracket \mathbf{p} \varphi \rrbracket_M \Leftrightarrow s \in \llbracket \mathbf{p} \varphi[x/y] \rrbracket_M.$$

PRINCIPLE 7 (*Preservation of Identities.*) Preferential readings should be insensitive to substitutions of equals:

$$\forall t_1, t_2 \in \mathbf{Var} \cup \mathbf{Con} : s \in \llbracket \mathbf{p} \varphi \wedge t_1 = t_2 \rrbracket_M \Leftrightarrow s \in \llbracket \mathbf{p} \varphi[t_1/t_2] \rrbracket_M.$$

4.4 Preferential Dynamic Disambiguation of Pronouns

We will now account for the discourse examples with ambiguous pronouns discussed in Section 2 using the first-order preferential discourse logic defined here.

4.4.1 Single-preferential Structure

We will first examine the single-preferential structure of the ‘John met Bill’ sentences (1)–(4). Assume the single-preferential extensions \mathcal{M} of the models presented in Table 5 subject to Principles 1, 5, 6, and 7. This model, together with the background information Γ containing (27) and (28), yields the intended defeasible conclusions as follows:

$$(29) \quad \begin{array}{ll} x = \mathbf{j} \wedge y = \mathbf{b}, \text{Meet } xy, \text{Greet } uv & \approx_{\mathcal{M}_\Gamma}^{\min + \mu} (u = \mathbf{j} \wedge v = \mathbf{b}) \\ x = \mathbf{b} \wedge y = \mathbf{j}, \text{Meet } xy, \text{Greet } uv & \approx_{\mathcal{M}_\Gamma}^{\min + \mu} (u = \mathbf{b} \wedge v = \mathbf{j}). \end{array}$$

This class also entails the invalidity of this kind of a determinate resolution for the ‘John and Bill met’-case (3):

$$(30) \quad \begin{array}{ll} x = \mathbf{j} \wedge y = \mathbf{b}, \text{Meet } xy \wedge \text{Meet } yx, & \\ \text{Greet } uv & \not\approx_{\mathcal{M}_\Gamma}^{\min + \mu} (u = \mathbf{j} \wedge v = \mathbf{b}). \end{array}$$

The underlying reason is that the preferential meaning of $\text{Meet } xy \wedge \text{Meet } yx$ may be different from that of $\text{Meet } xy$ or $\text{Meet } yx$, although these three sentences all have the same indefeasible meaning in \mathcal{M}_Γ .

For discourse (1) extended with the sentence *John greeted back* in (4), the defeasible conclusion of the first discourse in (29) will be invalid over \mathcal{M}_Γ :

$$(31) \quad \begin{array}{ll} x = \mathbf{j} \wedge y = \mathbf{b}, \text{Meet } xy, \text{Greet } uv, & \\ \text{Greet } xu & \not\approx_{\mathcal{M}_\Gamma}^{\min + \mu} (u = \mathbf{j} \wedge v = \mathbf{b}). \end{array}$$

The reason is that for every model $M \in \mathcal{M}_\Gamma$:

$$(32) \quad \forall s \in S : s \in \langle\langle x = \mathbf{j} \wedge y = \mathbf{b}, \text{Meet } xy, \text{Greet } uv \rangle\rangle_{M, \min}^{+\mu} \Rightarrow \llbracket \text{Greet } xu \rrbracket_{M, s}^{+\mu} = \emptyset.$$

4.4.2 Double-preferential Structure

We will now illustrate how the overriding effects of commonsense preferences illustrated in (8) and (9) come about in a double-preferential extension of the DML-setting in Table 5. In these cases, we hypothesized that the commonsense preferences about hitting / injuring / breaking override the syntactic preferences underlying the ‘John met Bill’ examples (1)–(4). We postulate the following double-preferential background for the ‘hitting’ scene:

$$(33) \quad \begin{aligned} & [\mathbf{p}_1 \text{ Hit } xy]^{+\mu} [\mathbf{p}_1 \text{ Injured } v]^{+\mu} v = x \quad \text{and} \\ & [\mathbf{p}_2 \text{ Hit } xy]^{+\mu} [\text{Injured } v]^{+\mu} v = y \end{aligned}$$

The \mathbf{p}_2 -class is associated with commonsense preferences with a higher preferential rank, while the \mathbf{p}_1 -class is associated with ‘syntactic’ preferences with a lower preferential rank. Note that we take the commonsense impact of the word *Hit* so strongly that every *Injured v*-continuation — not only the preferred readings of this sentence — leads to the defeasible conclusion that the hittee is the one who must be injured.

The above double-preferential account also enables a formal distinction among discourses F (same as (8) involving Bill), G (involving Schwarzenegger), and H (involving the Terminator) in Table 1, whose differences are exhibited in the survey results presented in Table 2.

Let \mathcal{N} be the class of double-preferential enrichments of the models of Table 5 satisfying the same principles as \mathcal{M} for both preference classes. When Δ represents the set containing the two preferential update rules given in (33), we obtain a determinate preference for F:

$$(34) \quad x = \mathbf{j} \wedge y = \mathbf{b}, \text{Hit } xy, \text{Injured } v \models_{\mathcal{N}_\Delta}^{\min+\mu} v = \mathbf{b}.$$

Let Δ' be the extension of Δ enriched with the following additional commonsense rules, where \mathbf{sh} denotes Schwarzenegger:

$$(35) \quad [\mathbf{p}_2 \text{ Injured } x]^{+\mu} [x = \mathbf{sh}]^{+\mu} \perp.$$

This rule says that if something is injured, then it is not expected to be Schwarzenegger. We then obtain a case of indeterminacy for G:

$$(36) \quad \begin{aligned} & x = \mathbf{j} \wedge y = \mathbf{sh}, \text{Hit } xy, \text{Injured } v \not\models_{\mathcal{N}_{\Delta'}}^{\min+\mu} v = \mathbf{sh} \text{ and} \\ & x = \mathbf{j} \wedge y = \mathbf{sh}, \text{Hit } xy, \text{Injured } v \not\models_{\mathcal{N}_{\Delta'}}^{\min+\mu} v = \mathbf{j}. \end{aligned}$$

Let Δ'' be the union of Δ and the following additional rules, where the constant \mathbf{tm} denotes the Terminator:

$$(37) \quad [\mathbf{j} = \mathbf{tm}]^{+\mu} \perp \text{ and } [\text{Injured } \mathbf{tm}]^{+\mu} \perp.$$

The second sentence says that the Terminator cannot be injured. This background information establishes the preferred meaning of H:

$$(38) \quad x = \mathbf{j} \wedge y = \mathbf{tm}, \text{Hit } xy, \text{Injured } v \models_{\mathcal{N}_{\Delta''}}^{\min+\mu} v = \mathbf{j}.$$

Substitution of $\Theta = \Delta' \cup \Delta''$ for Δ in (34), for Δ' in (36) and for Δ'' in (38) yields the same conclusions as above. In summary, if Θ was our background knowledge, then the discourse F predicts that Bill is injured, while G yields indeterminacy in its preferential meaning. Discourse H preferentially entails that John is injured.

5 Conclusions and Future Prospects

As a general logical basis for an integrated model of discourse semantics and pragmatics, we have combined dynamics and preferential reasoning in a dynamic modal logic setting. This logical setting encodes the basic discourse pragmatic properties of dynamicity, indeterminacy, defeasibility, and preference class interactions posited in an earlier linguistic analysis of the preferential effects on ambiguous pronouns. It also provides a logical architecture in which to implement a set of meta-constraints that regulates the general interplay of defeasible and infeasible static and dynamic interpretation. We have given a number of such meta-constraint candidates here. Further logical and empirical investigations are needed before we can choose the exact set of constraints we need.

We demonstrated how a general model theory of dynamic logic can be enriched with a preferential structure to result in a relatively simple preferential model theory. We defined the preferential dynamic entailments over given pieces of discourse, which predict that preferential information is used as much as possible and as early as possible to conclude discourse interpretations. That is, earlier defeasible conclusions are harder to defeat than more recent ones. We have also defined a logical machinery for predicting overriding relationships among preference classes. Overriding takes place when later infeasible information defeats earlier preferential conclusions, or when a reading corresponding to a preference class of a higher priority becomes empty and a lower preference class takes over. These preference class overrides give rise to conflict resolutions that are not predictable from straightforward applications of the Penguin Principle.

Although our focus is on pronoun resolution preferences in this paper, we hope that our logical machinery is also adequate for characterizing the conflict resolution patterns among various preferences and preference classes relevant to a wider range of discourse phenomena. The present perspective of preference interactions assumes that preferences belong to different classes, or modules, and that there are certain common conflict resolution patterns *within* each class and *across* different classes. Class-internal preference interactions yield either determinate or indeterminate preferences. Class-external preference interactions are dictated by certain preexisting class-level overriding relations, according to which the conflicts among the respective conclusions coming from each preference class are either resolved (by class-level overrides), ending up with the preferential conclusions of the highest preference class (whether it is determinate or indeterminate), or unresolved, leading to mixed-class preferential ambiguities. We would like to investigate the applicability of this perspective to a wider range of discourse phenomena.²⁶

The present logical characterization of preferential dynamics may be extended and/or revised in two major areas. One is the application of actions other than updates, $+\mu$. For example, discourse-level repairs as in (5) also require reductions, $-$, and/or downdates, $-\mu$. The other is the relational definition of preferences on the basis of an additional structuring of the information order \sqsubseteq instead of the static interpretation function $\llbracket \cdot \rrbracket$. Such an alternative definition would enable us to implement ‘graded’ preferences (Delgrande,

²⁶It is encouraging that the recent spread of Optimality Theory from phonology (Prince and Smolensky, 1993) to syntax (e.g., MIT Workshop on Optimality in Syntax, 1995) seems to indicate the descriptive adequacy of the fundamental preference interaction scheme, where potentially conflicting defeasible conclusions compete for the ‘maximal harmony.’

1988), that is, every state gets a certain preferential status with respect to a proposition. In our paper, states were simply declared to be preferential or nonpreferential with respect to a proposition. Graded preferences may be required for fine-tuning and coordinating the overall discourse pragmatics. A question related to this topic is whether the use of graded preferences would make the setting of multiple preference classes superfluous.

We might also be able to extend the framework to cover on-line sentence processing pragmatics, where the word-by-word or constituent-by-constituent dynamicity affects the meaning of the utterance being interpreted. The utterance-internal garden path and repair phenomena will then be treated analogously to the discourse-level counterparts.

References

- [Asher and Lascarides1995] Asher, Nicholas and Alex Lascarides. 1995. Lexical disambiguation in a discourse context. *Journal of Semantics*, 12(1):69–108. Special Issue on Lexical Semantics, Part I.
- [Asher and Morreau1991] Asher, Nicholas and Michael Morreau. 1991. Commonsense entailment: A modal theory of non-monotonic reasoning. In J. van Eijck, editor, *Logics in AI / JELIA90*, volume 478 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, Heidelberg, pages 1–30.
- [Benthem1991] Benthem, Johan Van. 1991. *Language in Action*, volume 130 of *Studies in Logic and the Foundations of Mathematics*. North Holland, Amsterdam.
- [Benthem and Cepparello1994] Benthem, Johan Van and Giovanna Cepparello. 1994. Tarskian variations; dynamic parameters in classical logic. Technical Report CS-R9419, CWI, Amsterdam.
- [Cohen and Levesque1990] Cohen, Phil and Hector Levesque. 1990. Intention is choice with commitment. *Artificial Intelligence Journal*, 42:213–261.
- [de Rijke1992] de Rijke, Maarten. 1992. A system of dynamic modal logic. Technical Report 92-170, CSLI, Stanford, CA. to appear in the *Journal of Philosophical Logic*.
- [Delgrande1988] Delgrande, J. 1988. An approach to default reasoning based on first-order conditional logic. *Artificial Intelligence Journal*, 36:63–90.
- [Fitting1969] Fitting, Melvin. 1969. *Intuitionistic Logic: Model Theory and Forcing*. Studies in Logic and the Foundations of Mathematics. North Holland, Amsterdam.
- [Grice1975] Grice, H. Paul. 1975. Logic and conversation. In P. Cole and J. Morgan, editors, *Speech Acts: Syntax and Semantics*, volume 3. Academic Press, New York, pages 41–58.
- [Groenendijk, Stokhof, and Veltman1996] Groenendijk, J., M. Stokhof, and F. Veltman. 1996. Coreference and modality. In Shalom Lappin, editor, *The Handbook of Contemporary Semantic Theory*. Blackwell, Oxford, UK, pages 179–213.

- [Groenendijk and Stokhof1991] Groenendijk, Jeroen and Martin Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100.
- [Jaspars and Krahmer1996] Jaspars, Jan and Emiel Krahmer. 1996. A programme of modal unification of dynamic theories. In P. Dekker and M. Stokhof, editors, *Proceedings of the Tenth Amsterdam Colloquium*. ILLC, Amsterdam, pages 425–444.
- [Kameyama1996] Kameyama, Megumi. 1996. Indefeasible semantics and defeasible pragmatics. In M. Kanazawa, C. Piñon, and H. de Swart, editors, *Quantifiers, Deduction, and Context*. CSLI, Stanford, CA, pages 111–138.
- [Kehler1995] Kehler, Andrew. 1995. *Interpreting Cohesive Forms in the Context of Discourse Inference*. Ph.D. thesis, Harvard University, Cambridge, MA, June. TR-11-95, Center for Research in Computing Technology.
- [Lascarides and Asher1993] Lascarides, Alex and Nicholas Asher. 1993. Temporal interpretation, discourse relations, and commonsense entailment. *Linguistics and Philosophy*, 16:437–493.
- [Levinson1983] Levinson, Stephen C. 1983. *Pragmatics*. Cambridge Textbooks in Linguistics. Cambridge University Press, Cambridge, U.K.
- [McCarthy and Hayes1969] McCarthy, John and Patrick Hayes. 1969. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence*, volume 4. Edinburgh University Press, Edinburgh, pages 463–502.
- [Muskens, Benthem, and Visser1997] Muskens, Reinhard, Johan Van Benthem, and Albert Visser. 1997. Dynamics. In Johan Van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*. Elsevier Science, pages 587–648.
- [Prince and Smolensky1993] Prince, A. and P. Smolensky. 1993. Optimality theory: Constraint interaction in generative grammar. Technical Report 2, Center for Cognitive Science, Rutgers University, New Brunswick, NJ.
- [Reiter1980] Reiter, Raymond. 1980. A logic for default reasoning. *Artificial Intelligence Journal*, 13:81–132.
- [Shoham1988] Shoham, Yoav. 1988. *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence*. The MIT Press Series in Artificial Intelligence. MIT Press, Cambridge, MA.
- [Troelstra and Van Dalen1988] Troelstra, Anne and Dirk Van Dalen. 1988. *Constructivism in Mathematics, volume I*. Studies in Logic and the Foundations of Mathematics. North Holland, Amsterdam.
- [Veltman1991] Veltman, Frank. 1991. Defaults in update semantics. Technical Report LP-91-02, Department of Philosophy, University of Amsterdam. To appear in the *Journal of Philosophical Logic*.